arXiv:1602.02172v1 [cs.LG] 5 Feb 2016

On Column Selection in Approximate Kernel Canonical Correlation Analysis

Weiran Wang
Toyota Technological Institute at Chicago
6045 S Kenwood Ave., Chicago IL 60637
weiranwang@ttic.edu

February 9, 2016

Abstract

We study the problem of column selection in large-scale kernel canonical correlation analysis (KCCA) using the Nyström approximation, where one approximates two positive semi-definite kernel matrices using "landmark" points from the training set. When building low-rank kernel approximations in KCCA, previous work mostly samples the landmarks uniformly at random from the training set. We propose novel strategies for sampling the landmarks non-uniformly based on a version of statistical leverage scores recently developed for kernel ridge regression. We study the approximation accuracy of the proposed non-uniform sampling strategy, develop an incremental algorithm that explores the path of approximation ranks and facilitates efficient model selection, and derive the kernel stability of out-of-sample mapping for our method. Experimental results on both synthetic and real-world datasets demonstrate the promise of our method.

1 Introduction

Kernel canonical correlation analysis (KCCA, Lai and Fyfe, 2000, Akaho, 2001, Melzer et al., 2001, Bach and Jordan, 2002, Hardoon et al., 2004) is the kernel extension of the classic canonical correlation analysis (CCA, Hotelling, 1936) algorithm, and has found wide application for analyzing multi-view data in different areas, such as computer vision [Socher and Li, 2010], speech recognition [Rudzicz, 2010, Arora and Livescu, 2012], natural language processing [Vinokourov et al., 2003, Li and Shawe-Taylor, 2005, Hardoon et al., 2007, Hodosh et al., 2013], computational biology [Yamanishi et al., 2004, Hardoon et al., 2007, Blaschko et al., 2011], and statistics [Bach and Jordan, 2002, Fukumizu et al., 2007, Lopez-Paz et al., 2013].

The main idea of CCA is to *linearly* project two random vectors into a lower-dimensional subspace so that the projections are maximally correlated. To extend CCA to nonlinear mappings with greater representation power, KCCA first maps the input observations of each view into Reproducing Kernel Hilbert Spaces (RKHS) and then applies linear CCA in the RKHS. Formally, given a dataset of N pairs of observations $(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_N, \mathbf{y}_N)$ of the random vectors $(\mathbf{x}, \mathbf{y})$, where $\mathbf{x}_i \in \mathbb{R}^{d_x}$ and $\mathbf{y}_i \in \mathbb{R}^{d_y}$ for $i = 1, \ldots, N$, KCCA maps the inputs into $\phi(\mathbf{x}_i) = k_x(\mathbf{x}_i, \cdot) \in \mathcal{H}_x$ and $\psi(\mathbf{y}_i) = k_y(\mathbf{y}_i, \cdot) \in \mathcal{H}_y$, where k_x and k_y are positive definite kernels associated with RKHS spaces $\mathcal{H}_x$ and $\mathcal{H}_y$ respectively, and then solves the following *regularized* objective

$$\begin{aligned} &\max_{f \in \mathcal{H}_x, g \in \mathcal{H}_y} \langle f, \mathbf{\Sigma}_{xy} g \rangle_{\mathcal{H}_x} \\ \text{s.t.} \quad &\langle f, \mathbf{\Sigma}_{xx} f \rangle_{\mathcal{H}_x} + 2\lambda_1 \|f\|^2_{\mathcal{H}_x} = 1 \\ &\langle g, \mathbf{\Sigma}_{yy} g \rangle_{\mathcal{H}_y} + 2\lambda_2 \|g\|^2_{\mathcal{H}_y} = 1 \end{aligned} \tag{1}$$

where $(\lambda_1, \lambda_2) > 0$ are regularization parameters that help avoid trivial solutions (and enable matrix inversions), and

the cross- and auto- covariance operators are defined as

$$\langle f, \mathbf{\Sigma}_{xy} g \rangle_{\mathcal{H}_x} = \frac{1}{N} \sum_{i=1}^{N} \left\langle f, \bar{\phi}(\mathbf{x}_i) \right\rangle_{\mathcal{H}_x} \left\langle g, \bar{\psi}(\mathbf{y}_i) \right\rangle_{\mathcal{H}_y},$$
$$\langle f, \mathbf{\Sigma}_{xx} f \rangle_{\mathcal{H}_x} = \frac{1}{N} \sum_{i=1}^{N} \left\langle f, \bar{\phi}(\mathbf{x}_i) \right\rangle_{\mathcal{H}_x} \left\langle f, \bar{\phi}(\mathbf{x}_i) \right\rangle_{\mathcal{H}_x},$$

with $\bar{\phi}(\mathbf{x}_i) = \phi(\mathbf{x}_i) - \frac{1}{N}\sum_{j=1}^{N} \phi(\mathbf{x}_j)$ and $\bar{\psi}(\mathbf{y}_i) = \psi(\mathbf{y}_i) - \frac{1}{N}\sum_{j=1}^{N} \psi(\mathbf{y}_j)$. From now on, we denote by $\mathbf{K}_1$ and $\mathbf{K}_2$ the *uncentered* kernel (Gram) matrices, i.e., $(\mathbf{K}_1)_{ij} = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle$ and $(\mathbf{K}_2)_{ij} = \langle \psi(\mathbf{y}_i), \psi(\mathbf{y}_j) \rangle$, and denote by $\overline{\mathbf{K}}_1 = \mathbf{H}\mathbf{K}_1\mathbf{H}$ and $\overline{\mathbf{K}}_2 = \mathbf{H}\mathbf{K}_2\mathbf{H}$ their centered version, where $\mathbf{H} = \mathbf{I} - \frac{1}{N}\mathbf{1}\mathbf{1}^\top \in \mathbb{R}^{N \times N}$.

Following the derivation of Bach and Jordan [2002], it suffices to consider f and g that lie in the span of $\bar{\phi}(\mathbf{x}_i)$ and $\bar{\psi}(\mathbf{y}_i)$ respectively, i.e., $f = \sum_{i=1}^{N} \alpha_i \bar{\phi}(\mathbf{x}_i)$ and $g = \sum_{i=1}^{N} \beta_i \bar{\psi}(\mathbf{y}_i)$, and the optimal coefficients $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_N]^\top$, $\boldsymbol{\beta} = [\beta_1, \ldots, \beta_N]^\top \in \mathbb{R}^N$ satisfy the following eigenvalue system[1]

$$\begin{bmatrix} \mathbf{0} & \mathbf{T} \\ \mathbf{T}^\top & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha}' \\ \boldsymbol{\beta}' \end{bmatrix} = \rho \begin{bmatrix} \boldsymbol{\alpha}' \\ \boldsymbol{\beta}' \end{bmatrix}, \tag{2}$$

where $\boldsymbol{\alpha}' = \frac{1}{\sqrt{N}}(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})\boldsymbol{\alpha}$, $\boldsymbol{\beta}' = \frac{1}{\sqrt{N}}(\overline{\mathbf{K}}_2 + N\lambda_2\mathbf{I})\boldsymbol{\beta}$ with $\|\boldsymbol{\alpha}'\| = \|\boldsymbol{\beta}'\| = 1$, and

$$\mathbf{T} = (\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{K}}_1\overline{\mathbf{K}}_2(\overline{\mathbf{K}}_2 + N\lambda_2\mathbf{I})^{-1}, \tag{3}$$

and ρ is the largest eigenvalue of the system and the optimal objective of (1). Alternatively, ρ is the largest singular value of $\mathbf{T}$, while $\boldsymbol{\alpha}'$ and $\boldsymbol{\beta}'$ are the corresponding left and right singular vectors. The view 1 projection mapping is

$$f(\mathbf{x}) = \sum_{i=1}^{N} \alpha_i \left\langle \bar{\phi}(\mathbf{x}_i), \bar{\phi}(\mathbf{x}) \right\rangle \quad \in \mathbb{R}. \tag{4}$$

L-dimensional projections can be obtained by simultaneously extracting the top L eigenvectors of (2).

Although KCCA has a closed-form solution, for applications where the training set size N is large, it is computationally infeasible to evaluate and store the $N \times N$ kernel matrices, let alone compute $\mathbf{T}$ and its rank-L singular value decomposition (SVD) which are of time complexity $\mathcal{O}(N^3)$ and $\mathcal{O}(LN^2)$ respectively. As a result, various low-rank kernel approximation techniques have been proposed to alleviate these issues, including incomplete Cholesky decomposition [Bach and Jordan, 2002], partial Gram-Schmidt [Hardoon et al., 2004], incremental SVD [Arora and Livescu, 2012], and random Fourier features [Lopez-Paz et al., 2014]. Another popular low-rank kernel approximation method is the Nyström approximation [Williams and Seeger, 2001, Fowlkes et al., 2004, Kumar et al., 2012, Gittens and Mahoney, 2013]: By sampling a set of M landmark points from the training set with index set $I = \{i_1, \ldots, i_M\}$, a symmetric positive semi-definite (PSD) matrix $\mathbf{K} \in \mathbb{R}^{N \times N}$ is approximated as

$$\mathbf{K} \approx \mathbf{L} := \mathbf{C}\mathbf{W}^\dagger\mathbf{C}^\top \tag{5}$$

where $\mathbf{C} \in \mathbb{R}^{N \times M}$ contains the sampled columns of $\mathbf{K}$ indexed by I, and $\mathbf{W} \in \mathbb{R}^{M \times M}$ is the square sub-matrix of $\mathbf{K}$ with rows and columns both indexed by I.

One can build Nyström approximations $\mathbf{L}_1$ of rank M_1 for $\mathbf{K}_1$, and $\mathbf{L}_2$ of rank M_2 for $\mathbf{K}_2$ in KCCA; we denote the resulting algorithm NKCCA. NKCCA typically outperforms approximate KCCA using random Fourier features [Lopez-Paz et al., 2014] for the same rank, even when a naive uniform sampling strategy is adopted.[2]

Our contributions First, we propose a non-uniform sampling strategy for NKCCA, and prove its approximation guarantee in terms of the kernel canonical correlation. We show that, by carefully selecting small fractions of columns of $\mathbf{K}_1$ and $\mathbf{K}_2$ in Nyström approximations, we obtain an NKCCA solution that is an accurate estimate of the solution to exact KCCA, and that to achieve the same accuracy, the proposed strategy may require fewer columns to be sampled than does uniform sampling. This is the first approximation guarantee for NKCCA that we are aware of, despite its

[1] Note that Bach and Jordan [2002] used approximations of the form $\langle f, \mathbf{\Sigma}_{xx} f \rangle_{\mathcal{H}_x} + 2\lambda_1 \|f\|^2_{\mathcal{H}_x} = \frac{1}{N}\boldsymbol{\alpha}^\top\overline{\mathbf{K}}_1^2\boldsymbol{\alpha} + 2\lambda_1\boldsymbol{\alpha}^\top\overline{\mathbf{K}}_1\boldsymbol{\alpha} \approx \frac{1}{N}\boldsymbol{\alpha}^\top(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^2\boldsymbol{\alpha}$, which still lead to a consistent estimate of the kernel canonical correlation.

[2] This work uses $\mathbf{C}(\mathbf{W}^\dagger)^{1/2} \in \mathbb{R}^{N \times M}$ as the new data matrix in each view and run linear CCA on top. One can show this is equivalent to solving an eigenvalue system similar to (2) with Nyström kernel approximations, using the argument in Lopez-Paz et al. [2014, Sec. 4].

common use in practice [Lopez-Paz et al., 2014, Xie et al., 2015, Wang et al., 2015]. Our strategy is motivated by that of Alaoui and Mahoney [2015] for kernel ridge regression, who recently showed that if the columns are sampled according to a distribution depending on a version of statistical leverage scores, the statistical performance can be guaranteed when low-rank kernel approximations are used. Notice that CCA is closely related to regression, as its objective is to find a common subspace where the two views are maximally correlated and predictive of each other. Second, motivated by the algorithm of Rudi et al. [2015] for kernel ridge regression, we propose an incremental Nyström approximation algorithm for NKCCA, which explores the entire path of approximation ranks and facilitates efficient model selection. While the first two points are motivated by recent work, the column selection problem here is more interesting and challenging since there are two kernel matrices in KCCA. Third, we derive the kernel stability of NKCCA, i.e., the perturbation of the projection mapping due to kernel approximations used in training, and our approach can be extended to other approximate KCCA algorithms with spectral error bounds in $\mathbf{T}$. As far as we know, this is the first kernel stability result for KCCA.

2 Column selection in KCCA

Notations Boldface capital letters denote matrices; boldface lower-case letters denote column vectors; without boldface, lower-case letters denote scalars. $\mathbf{1}$ denotes the column vector of all ones, and $\mathbf{I}$ denotes the identity matrix with appropriate dimensions. $[\mathbf{u}; \mathbf{v}]$ denotes the concatenation of two vectors $\mathbf{u}$ and $\mathbf{v}$. For any matrix $\mathbf{A}$, we use $\sigma_j(\mathbf{A})$ to denote its j-th largest singular value, $\|\mathbf{A}\| = \sigma_1(\mathbf{A})$ its spectral norm, and $\mathbf{A}^\dagger$ its pseudo-inverse. We use $\preceq$ to indicate PSD ordering, i.e., for two PSD matrices $\mathbf{A}$ and $\mathbf{B}$, $\mathbf{A} \preceq \mathbf{B}$ if and only if $\mathbf{B} - \mathbf{A}$ is PSD.

We consider the following rank-M_1 approximation to the *uncentered* view 1 kernel matrix

$$\mathbf{K}_1 \approx \mathbf{L}_{\gamma_1} := \mathbf{K}_1\mathbf{S}_1(\mathbf{S}_1^\top\mathbf{K}_1\mathbf{S}_1 + N\gamma_1\mathbf{I})^\dagger\mathbf{S}_1^\top\mathbf{K}_1, \tag{6}$$

where $\gamma_1 \geq 0$, and $\mathbf{S}_1 \in \mathbb{R}^{N\times M_1}$ is the sampling matrix such that $(\mathbf{S}_1)_{im} > 0$ if column i is chosen at the m-th trial of sampling the M_1 columns of $\mathbf{K}_1$, and zero otherwise. The centered kernel matrix is then approximated as $\overline{\mathbf{K}}_1 \approx \overline{\mathbf{L}}_{\gamma_1} := \mathbf{H}\mathbf{L}_{\gamma_1}\mathbf{H}$. For the view 2 kernel $\mathbf{K}_2$, we have the rank M_2, $\mathbf{S}_2 \in \mathbb{R}^{N\times M_2}$, $\mathbf{L}_{\gamma_2}$, and $\overline{\mathbf{L}}_{\gamma_2}$ defined analogously. Notice that (6) is a slight generalization of the standard Nyström approximation; by setting $\gamma_1 = \gamma_2 = 0$ and the nonzero entries of $\mathbf{S}_1$ and $\mathbf{S}_2$ to 1, we recover the scheme in (5). We denote the version of $\mathbf{L}_{\gamma_1}$ (resp. $\overline{\mathbf{L}}_{\gamma_1}$) with $\gamma_1 = 0$ as $\mathbf{L}_1$ (resp. $\overline{\mathbf{L}}_1$), and similarly for view 2. We omit subscripts $1/2$ below if the result holds for both views.

Alaoui and Mahoney [2015, Lemma 1] provide the following basic deterministic characterizations of $\mathbf{L}_\gamma$.

Lemma 1. *Let $\mathbf{K} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top$ where $\mathbf{U}$ is orthogonal and $\mathbf{\Sigma}$ diagonal non-negative. For $\gamma > 0$, we have*

$$\mathbf{L}_\gamma \preceq \mathbf{L} \preceq \mathbf{K}. \tag{7}$$

Moreover, let

$$\mathbf{D} = \mathbf{\Phi} - \mathbf{\Phi}^{1/2}\mathbf{U}^\top\mathbf{S}\mathbf{S}^\top\mathbf{U}\mathbf{\Phi}^{1/2} \tag{8}$$

with $\mathbf{\Phi} = \mathbf{\Sigma}(\mathbf{\Sigma} + N\gamma\mathbf{I})^{-1}$. If $\|\mathbf{D}\| \leq t$ for $t \in (0, 1)$, then

$$\mathbf{0} \preceq \mathbf{K} - \mathbf{L}_\gamma \preceq \frac{N\gamma}{1-t}\mathbf{I}. \tag{9}$$

Based on Lemma 1, we obtain the following key lemma (see its proof in the supplementary materials).

Lemma 2. *Let $\gamma > 0$, and assume the conditions of Lemma 1 hold. Then we have*

$$\max\left\{\left\|\mathbf{K}(\mathbf{K} + N\lambda\mathbf{I})^{-1} - \mathbf{L}(\mathbf{L} + N\lambda\mathbf{I})^{-1}\right\|,\ \left\|\mathbf{K}(\mathbf{K} + N\lambda\mathbf{I})^{-1} - \mathbf{L}_\gamma(\mathbf{L}_\gamma + N\lambda\mathbf{I})^{-1}\right\|\right\} \leq \frac{\gamma/\lambda}{1-t}.$$

Moreover, the same bound also holds if $(\mathbf{K}, \mathbf{L}_\gamma, \mathbf{L})$ are replaced by $(\overline{\mathbf{K}}, \overline{\mathbf{L}}_\gamma, \overline{\mathbf{L}})$ simultaneously.

Guaranteeing that the conditions in the lemma hold boils down to controlling the spectral norm of $\mathbf{D}$ defined in (8). Observe that $\mathbf{D}$ is the error of approximating $\mathbf{\Phi} = \mathbf{\Sigma}(\mathbf{\Sigma}+\gamma\mathbf{I})^{-1}$ or equivalently the multiplication $(\mathbf{U}\mathbf{\Phi}^{1/2})^\top \cdot (\mathbf{U}\mathbf{\Phi}^{1/2})$ by a subset of the rows of $\mathbf{U}\mathbf{\Phi}^{1/2}$ (indexed by nonzero entries of $\mathbf{S}$), and the optimal strategy (in Frobenius

norm) is to sample each row with a probability proportional to its squared length [Drineas et al., 2006]. This motivates the following definition of γ-*ridge leverage scores* [Alaoui and Mahoney, 2015]: For $\gamma > 0$, the γ-ridge leverage scores associated with $\mathbf{K}$ and the parameter γ are

$$l_i(\mathbf{K}, \gamma) = \sum_{j=1}^{N} \frac{\sigma_j(\mathbf{K})}{\sigma_j(\mathbf{K}) + N\gamma} \mathbf{U}_{ij}^2, \quad j = 1, \ldots, N. \tag{10}$$

Notice that $l_i(\mathbf{K}, \gamma)$ is precisely the squared norm of the i-th row of $\mathbf{U}\mathbf{\Phi}^{1/2}$, and equivalently $\left(\mathbf{K}(\mathbf{K} + N\gamma\mathbf{I})^{-1}\right)_{ii}$. Moreover, the *effective dimensionality* of $\mathbf{K}$ with parameter γ is defined as the sum of γ-ridge leverage scores:

$$d_{\text{eff}}(\mathbf{K}, \gamma) = \sum_{i=1}^{N} l_i(\mathbf{K}, \gamma) = \operatorname{tr}\left(\mathbf{K}(\mathbf{K} + N\gamma\mathbf{I})^{-1}\right). \tag{11}$$

The following lemma provides the number of rows that need to be sampled for $\|\mathbf{D}\|$ to be small. It is essentially an application of Theorem 2 of Alaoui and Mahoney [2015] to the matrix multiplication $(\mathbf{U}\mathbf{\Phi}^{1/2})^\top \cdot (\mathbf{U}\mathbf{\Phi}^{1/2})$.

Lemma 3. *Let $M \leq N$ and $I = \{i_1, \ldots, i_M\}$ be a subset of $\{1, \ldots, N\}$ formed by M elements chosen randomly with replacement, according to the distribution*

$$Pr\,(choose\ i) = p_i \geq \beta \frac{l_i(\mathbf{K}, \gamma)}{d_{\mathit{eff}}(\mathbf{K}, \gamma)},\ i = 1, \ldots, N, \tag{12}$$

for some $\beta \in (0, 1]$. Let $\mathbf{S}$ be the corresponding sampling matrix such that $\mathbf{S}_{i_j j} = 1/\sqrt{M p_{i_j}}$ for $j = 1, \ldots, M$ and zero otherwise. If $M \geq 2\left(\frac{d_{\mathit{eff}}(\mathbf{K},\gamma)}{\beta t^2} + \frac{1}{3t}\right)\log\left(\frac{N}{\delta}\right)$, then with probability at least $1 - \delta$, we have $\|\mathbf{D}\| < t$.

Remark Lemma 3 allows the sampling distribution to be different from the optimal one by a factor of β, at the cost of slight over-sampling by a factor of roughly $1/\beta$ (assuming the first term in the lower bound of M is dominant). If the γ-ridge leverage scores are very non-uniform but we insist on using the uniform sampling strategy ($p_1 = \cdots = p_N = \frac{1}{N}$), the over-sampling rate is roughly $\frac{N l_1(\mathbf{K},\gamma)}{d_{\text{eff}}(\mathbf{K},\gamma)}$, which can be very large if $\mathbf{K}$ has a fast decaying spectrum. This shows the advantage of the more data-dependent sampling strategies. Alaoui and Mahoney [2015] also provide an efficient procedure for estimating γ-ridge leverage scores.

2.1 Approximation error of NKCCA

Our goal is to quantify the perturbation in the kernel canonical correlation ρ when $\mathbf{K}_1$ and $\mathbf{K}_2$ are replaced by $\mathbf{L}_1$ and $\mathbf{L}_2$ respectively. Now define

$$\tilde{\mathbf{T}} = (\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{L}}_1\overline{\mathbf{L}}_2(\overline{\mathbf{L}}_2 + N\lambda_2\mathbf{I})^{-1}, \tag{13}$$

and denote by $\tilde{\rho}$ its largest singular value with corresponding left and right singular vectors $\tilde{\boldsymbol{\alpha}}'$ and $\tilde{\boldsymbol{\beta}}'$ respectively. According to Weyl's inequality [Horn and Johnson, 1986], we have

$$\begin{aligned}
|\rho - \tilde{\rho}| &\leq \left\|\mathbf{T} - \tilde{\mathbf{T}}\right\| \\
&\leq \left\|(\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{L}}_1\overline{\mathbf{L}}_2(\overline{\mathbf{L}}_2 + N\lambda_2\mathbf{I})^{-1} - (\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{K}}_1\overline{\mathbf{L}}_2(\overline{\mathbf{L}}_2 + N\lambda_2\mathbf{I})^{-1}\right\| \\
&\quad + \left\|(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{K}}_1\overline{\mathbf{L}}_2(\overline{\mathbf{L}}_2 + N\lambda_2\mathbf{I})^{-1} - (\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{K}}_1\overline{\mathbf{K}}_2(\overline{\mathbf{K}}_2 + N\lambda_2\mathbf{I})^{-1}\right\| \\
&\leq \left\|(\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{L}}_1 - (\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{K}}_1\right\| + \left\|\overline{\mathbf{L}}_2(\overline{\mathbf{L}}_2 + N\lambda_2\mathbf{I})^{-1} - \overline{\mathbf{K}}_2(\overline{\mathbf{K}}_2 + N\lambda_2\mathbf{I})^{-1}\right\|
\end{aligned} \tag{14}$$

where we have used the triangle inequality, and the fact that $\left\|\overline{\mathbf{K}}(\overline{\mathbf{K}} + N\lambda\mathbf{I})^{-1}\right\| \leq 1$ in the two inequalities. It is then straightforward to bound each of the two terms using Lemma 2 and obtain the following guarantee.

Theorem 1. *Assume that, for constructing Nyström approximations $\mathbf{L}_1$ and $\mathbf{L}_2$, we sample M_1 columns from $\mathbf{K}_1$ and M_2 columns from $\mathbf{K}_2$ according to the distributions in* (12) *using ridge leverage scores $\{l_i(\mathbf{K}_1, \epsilon\lambda_1(1 - t_1)/2)\}$ and $\{l_i(\mathbf{K}_2, \epsilon\lambda_2(1 - t_2)/2)\}$ for some $t_1, t_2 \in [0, 1)$ respectively. If $M_1 \geq 2\left(\frac{d_{\mathit{eff}}(\mathbf{K}_1,\epsilon\lambda_1(1-t_1)/2)}{\beta t_1^2} + \frac{1}{3t_1}\right)\log\left(\frac{2N}{\delta}\right)$ and $M_2 \geq 2\left(\frac{d_{\mathit{eff}}(\mathbf{K}_2,\epsilon\lambda_2(1-t_2)/2)}{\beta t_2^2} + \frac{1}{3t_2}\right)\log\left(\frac{2N}{\delta}\right)$, then with probability at least $1 - \delta$, we have $|\rho - \tilde{\rho}| \leq \epsilon$.*

Proof. Setting $\gamma_1 = \epsilon\lambda_1(1-t_1)/2$ in Lemma 2 and using the stated M_1 in Lemma 3, we obtain that with probability at least $1-\delta/2$,

$$\left\|(\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{L}}_1 - (\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{K}}_1\right\| \leq \epsilon/2.$$

The same result can be obtained for view 2 and the theorem follows from (14) and a union bound. □

Remarks 1. It is important to note that, although our analysis largely depends on $\mathbf{L}_\gamma$ with $\gamma > 0$ (e.g., Lemma 1 & 2), it is not used in the actual algorithm. In fact, in Theorem 1, we only need a suitable value of γ to compute the estimated γ-ridge leverage scores (N numbers) and define the sampling distribution for each view.[3]

2. The theorem covers both non-uniform sampling and uniform sampling strategies, through the over-sampling factor β (see the remark after Lemma 3). In practice, we find that uniform sampling tends to work quite well already and ridge leverage scores may bring (moderate) further improvement.

3. There exist interesting trade-offs between accuracy and computation in the algorithm.

- From the definitions (10) and (11), we observe that $\{l_i(\mathbf{K},\gamma)\}$ and $d_{\text{eff}}(\mathbf{K},\gamma)$ are decreasing functions of γ. As a sanity check, when we require higher approximation accuracy (smaller ϵ), the d_{eff} terms in the lower bounds of M_1 and M_2 become larger, indicating that we need to sample more columns.
- It is easy to show that $d_{\text{eff}}(\mathbf{K},\gamma) \leq \operatorname{tr}(\mathbf{K})/N\gamma$. Therefore, for fixed t_1 and t_2, the necessary ranks M_1 and M_2 scale as $\mathcal{O}(1/\epsilon)$. Theoretically, this asymptotic dependence is much better than those obtained for approximate KCCA using data-independent random Fourier features, where ranks of kernel approximations (number of random features) scale as $\mathcal{O}(1/\epsilon^2)$ [Lopez-Paz et al., 2014, Theorem 4][4].

2.2 Incremental column selection

In practice, it is computationally expensive to compute exact ridge leverage scores. We can use the efficient algorithm of Alaoui and Mahoney [2015, Sec. 3.5] to compute their approximate values. This would introduce some approximation factor $\beta < 1$, and makes it difficult to determine the values of M_1 and M_2 a priori. Furthermore, in machine learning problems, we also care about the generalization ability of an algorithm, and large ranks that perform well on training data may not be optimal for test data.

These considerations motivate us to derive an incremental approach where we gradually increase M_1 and M_2 for the kernel approximations by sampling more columns according to the estimated ridge leverage scores, compute solutions along the path of different ranks, and monitor their performance on validation data to avoid over-fitting.

Note that $\bar{\mathbf{T}}$ is the multiplication of two matrices, each of the form $(\overline{\mathbf{L}}+N\lambda\mathbf{I})^{-1}\overline{\mathbf{L}}$. By the matrix inversion lemma,

$$\begin{aligned}
(\overline{\mathbf{L}} + N\lambda\mathbf{I})^{-1}\overline{\mathbf{L}} &= \mathbf{I} - N\lambda(\overline{\mathbf{L}} + N\lambda\mathbf{I})^{-1} \\
&= \mathbf{I} - \left(\mathbf{I} + \frac{1}{N\lambda}\mathbf{HKS}(\mathbf{S}^\top\mathbf{KS})^{-1}\mathbf{S}^\top\mathbf{KH}\right)^{-1} \\
&= \mathbf{I} - \left(\mathbf{I} - \mathbf{HKS}\left(N\lambda\mathbf{S}^\top\mathbf{KS} + \mathbf{S}^\top\mathbf{KHKS}\right)^{-1}\mathbf{S}^\top\mathbf{KH}\right) \\
&= \mathbf{HKS}\left(N\lambda\mathbf{S}^\top\mathbf{KS} + \mathbf{S}^\top\mathbf{KHKS}\right)^{-1}\mathbf{S}^\top\mathbf{KH}. \qquad (15)
\end{aligned}$$

Therefore, efficient incremental column selection hinges on inexpensive incremental updates of (15), where the time complexity mainly comes from computing the inverse.

When we add one more column in the Nyström approximation, the dimension of $\mathbf{S}^\top\mathbf{KS}$ and $\mathbf{S}^\top\mathbf{KHKS}$ both increase by 1. An efficient way of computing the inverses at each subsampling level is by exploiting rank-one Cholesky updates [Golub and van Loan, 1996]. We give the procedure for computing $\left(N\lambda\mathbf{S}^\top\mathbf{KS} + \mathbf{S}^\top\mathbf{KHKS}\right)^{-1}$ in Algorithm 1, where `cholupdate` is the rank-one update procedure implemented in many standard linear algebra libraries. The algorithm can output the Cholesky decompositions at any intermediate rank and the inverse is computed by efficiently solving two upper/lower triangular systems. A similar algorithm without the centering and scaling operations was used by Rudi et al. [2015] for incremental Nyström approximation in kernel ridge regression.

[3] The algorithm of Alaoui and Mahoney [2015] for estimating the scores does, however, construct an approximation of $\mathbf{L}_\gamma$.

[4] Although their bound is for expected spectral norm error, one can instead obtain a high probability bound when invoking the matrix Bernstein inequality.

Algorithm 1 Incremental algorithm for computing $\left(N\lambda\mathbf{S}^\top\mathbf{K}\mathbf{S}+\mathbf{S}^\top\mathbf{K}\mathbf{H}\mathbf{K}\mathbf{S}\right)^{-1}$ in Nyström approximation.

Input: Index set $I=\{i_1,\ldots,i_M\}$, regularization λ.

$s_1 \leftarrow \mathbf{S}_{i_1 1}$, $\mathbf{a}_1 \leftarrow s_1\mathbf{H}[\mathbf{K}_{i_1 1},\ldots,\mathbf{K}_{i_1 N}]^\top$, $\mathbf{A}_1 \leftarrow \mathbf{a}_1$

$d_1 \leftarrow \mathbf{a}_1^\top\mathbf{a}_1 + N\lambda s_1^2\mathbf{K}_{i_1 i_1}$, $\mathbf{R}_1 \leftarrow \sqrt{d_1}$.

for $m=2,\ldots,M$ **do**

$s_m \leftarrow \mathbf{S}_{i_m m}$

$\mathbf{a}_m \leftarrow s_m\mathbf{H}[\mathbf{K}_{i_m 1},\ldots,\mathbf{K}_{i_m N}]^\top$

$\mathbf{A}_m \leftarrow [\mathbf{A}_{m-1},\mathbf{a}_m]$

$\mathbf{b}_m \leftarrow s_m[s_1\mathbf{K}_{i_m i_1},\ldots,s_{m-1}\mathbf{K}_{i_m i_{m-1}}]^\top$

$\mathbf{c}_m \leftarrow \mathbf{A}_{m-1}^\top\mathbf{a}_m + N\lambda\mathbf{b}_m$

$d_m \leftarrow \mathbf{a}_m^\top\mathbf{a}_m + N\lambda s_m^2 K_{i_m i_m}$, $g_m \leftarrow \sqrt{1+d_m}$

$\mathbf{u}_m \leftarrow \left[\frac{\mathbf{c}_m}{1+g_m};\ g_m\right]$, $\mathbf{v}_m \leftarrow \left[\frac{\mathbf{c}_m}{1+g_m};\ -1\right]$

$\mathbf{R}_m \leftarrow \begin{bmatrix}\mathbf{R}_{m-1} & \mathbf{0}\\ \mathbf{0} & 0\end{bmatrix}$

$\mathbf{R}_m \leftarrow \texttt{cholupdate}(\mathbf{R}_m,\mathbf{u}_t,\text{'}+\text{'})$

$\mathbf{R}_m \leftarrow \texttt{cholupdate}(\mathbf{R}_m,\mathbf{v}_t,\text{'}-\text{'})$

end for

Output: $\left(N\lambda\mathbf{S}^\top\mathbf{K}\mathbf{S}+\mathbf{S}^\top\mathbf{K}\mathbf{H}\mathbf{K}\mathbf{S}\right)^{-1} = \mathbf{R}_M^{-1}(\mathbf{R}_M^\top)^{-1}$ where $\mathbf{R}_M \in \mathbb{R}^{M\times M}$ is upper triangular.

In NKCCA, our goal is to compute the SVD of $\tilde{\mathbf{T}} = (\overline{\mathbf{L}}_1 + N\lambda\mathbf{I})^{-1}\overline{\mathbf{L}}_1\overline{\mathbf{L}}_2(\overline{\mathbf{L}}_2 + N\lambda\mathbf{I})^{-1} \in \mathbb{R}^{N\times N}$. Thus we also incrementally compute the QR decompositions $\mathbf{H}\mathbf{K}_1\mathbf{S}_1 = \mathbf{Q}_1\mathbf{P}_1$ where $\mathbf{Q}_1 \in \mathbb{R}^{N\times M_1}$ has orthogonal columns and $\mathbf{P}_1 \in \mathbb{R}^{M_1\times M_1}$ is upper-triangular, and similarly $\mathbf{H}\mathbf{K}_2\mathbf{S}_2 = \mathbf{Q}_2\mathbf{P}_2$ using the modified Gram-Schmidt algorithm [Horn and Johnson, 1986]. This allows us to reduce the original SVD problem to computing the SVD of

$$\hat{\mathbf{T}} := \mathbf{P}_1\left(N\lambda\mathbf{S}_1^\top\mathbf{K}_1\mathbf{S}_1 + \mathbf{S}_1^\top\mathbf{K}_1\mathbf{H}\mathbf{K}_1\mathbf{S}_1\right)^{-1}\left(\mathbf{S}_1^\top\mathbf{K}_1\mathbf{H}\cdot\mathbf{H}\mathbf{K}_2\mathbf{S}_2\right)\left(N\lambda\mathbf{S}_2^\top\mathbf{K}_2\mathbf{S}_2 + \mathbf{S}_2^\top\mathbf{K}_2\mathbf{H}\mathbf{K}_2\mathbf{S}_2\right)^{-1}\mathbf{P}_2^\top$$

which is of much smaller dimensions ($\hat{\mathbf{T}} \in \mathbb{R}^{M_1\times M_2}$). By computing the SVD of $\hat{\mathbf{T}} = \hat{\mathbf{U}}\hat{\mathbf{\Sigma}}\hat{\mathbf{V}}^\top$, we recover the SVD of $\tilde{\mathbf{T}} = (\mathbf{Q}_1\hat{\mathbf{U}})\hat{\mathbf{\Sigma}}(\mathbf{Q}_2\hat{\mathbf{V}})^\top$. Then $\tilde{\boldsymbol{\alpha}}'$ (resp. $\tilde{\boldsymbol{\beta}}'$) corresponds to the first column of $\mathbf{Q}_1\hat{\mathbf{U}}$ (resp. $\mathbf{Q}_2\hat{\mathbf{V}}$).

We incrementally compute the "core matrix" $\left(\mathbf{S}_1^\top\mathbf{K}_1\mathbf{H}\cdot\mathbf{H}\mathbf{K}_2\mathbf{S}_2\right) \in \mathbb{R}^{M_1\times M_2}$ and save it in memory, and its dimensions grow by 1 each time we sample one more column. Assuming $M_1 = M_2 = M$, then the time complexity of computing the SVD of $\tilde{\mathbf{T}}$ as described above is $\mathcal{O}(M^3+M^2N)$. Several expensive steps, including forming the core matrix ($\mathcal{O}(M^2N)$), the incremental Cholesky decompositions ($\mathcal{O}(M^3)$), and the modified Gram-Schmidt algorithm ($\mathcal{O}(M^2N)$) have computations reused at different subsampling levels, resulting in significant time savings for model selection.

2.3 Out-of-sample mapping and kernel stability

After extracting $\tilde{\mathbf{T}}$'s left sigular vector $\tilde{\boldsymbol{\alpha}}'$ (of unit length), we compute the coefficients $\tilde{\boldsymbol{\alpha}} = \sqrt{N}(\overline{\mathbf{L}}_1 + N\lambda_1)^{-1}\tilde{\boldsymbol{\alpha}}'$, as approximation to $\boldsymbol{\alpha} = \sqrt{N}(\overline{\mathbf{K}}_1 + N\lambda_1)^{-1}\boldsymbol{\alpha}'$, for combining kernel affinities in the out-of-sample mapping. Since

$$(\overline{\mathbf{L}}_1 + N\lambda_1)^{-1} = \frac{1}{N\lambda_1}(\mathbf{I} - (\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\overline{\mathbf{L}}_1),$$

we can simply reuse the decomposition of (15) to compute the inverse. Furthermore, we can rewrite (4) as

$$\begin{aligned} f(\mathbf{x}) &= \sum\nolimits_{i=1}^N \alpha_i \left\langle \bar{\phi}(\mathbf{x}_i), \phi(\mathbf{x}) - \frac{1}{N}\sum\nolimits_{j=1}^N \phi(\mathbf{x}_j)\right\rangle \\ &= \sum\nolimits_{i=1}^N \alpha_i \left\langle \bar{\phi}(\mathbf{x}_i), \phi(\mathbf{x})\right\rangle + \text{const} \\ &= \mathbf{k}^\top\mathbf{H}\boldsymbol{\alpha} + \text{const} \end{aligned} \tag{16}$$

where the constant is independent of $\mathbf{x}$ (thus need not be computed), and $\mathbf{k} = [k_x(\mathbf{x}_1,\mathbf{x}),\ldots,k_x(\mathbf{x}_N,\mathbf{x})]^\top$.

Assuming the exact kernel function is used during testing, the out-of-sample mapping for NKCCA is (ignoring the same constant in (16))

$$\tilde{f}(\mathbf{x}) := \mathbf{k}^\top \mathbf{H} \tilde{\boldsymbol{\alpha}}. \tag{17}$$

We now study the perturbation in out-of-sample mapping resulting from the low-rank kernel approximations used in training NKCCA, referred to as "kernel stability" by Cortes et al. [2010]. The following theorem provides the kernel stability when there exists a non-zero singular value gap for $\mathbf{T}$ (used by exact KCCA). The condition on the kernel affinity being bounded, i.e., $k_x(\mathbf{x}, \mathbf{x}') \le c$ for all $\mathbf{x}$ and $\mathbf{x}'$, is verified with $c = 1$ for the Gaussian RBF kernel $k_x(\mathbf{x}, \mathbf{x}') = e^{-\frac{\|\mathbf{x}-\mathbf{x}'\|^2}{2\sigma^2}}$ for example.

Theorem 2. *Use the sampling strategy stated in Theorem 1 for NKCCA, so that* $\left\|\mathbf{T} - \tilde{\mathbf{T}}\right\| \le \epsilon$ *with probability at least* $1 - \delta$*. Furthermore, assume the following two conditions*

1. $\forall \mathbf{x}, \mathbf{x}'$*, we have* $k_x(\mathbf{x}, \mathbf{x}') \le c$*.*
2. *There exist* $r > 0$ *such that* $\sigma_1(\mathbf{T}) - \sigma_2(\mathbf{T}) \ge r$*, and* $\left\|\mathbf{T} - \tilde{\mathbf{T}}\right\| \le \frac{r}{2}$*.*

Then with the same probability, we have

$$\left| f(\mathbf{x}) - \tilde{f}(\mathbf{x}) \right| \le \left(\frac{1}{2} + \frac{4\sqrt{2}}{r} \right) \frac{c\epsilon}{\lambda_1} \qquad \forall \mathbf{x}.$$

Proof. The proof consists of three steps.

Step I. We first bound the perturbation in the unit singular vectors, i.e., $\left\|\boldsymbol{\alpha}' - \tilde{\boldsymbol{\alpha}}'\right\|$.

Let the complete set of left singular vectors of $\mathbf{T}$ be $\{\mathbf{u}_i\}_{i=1}^N$ (with $\mathbf{u}_1 = \boldsymbol{\alpha}'$), and the complete set of right singular vectors be $\{\mathbf{v}_i\}_{i=1}^N$ (with $\mathbf{v}_1 = \boldsymbol{\beta}'$), both of which constitute orthonormal basis sets of $\mathbb{R}^N$. Observe that the set of concatenated vectors $\{\mathbf{w}_i\}_{i=1}^{2N}$ with $\mathbf{w}_i = [\mathbf{u}_i; \mathbf{v}_i]$ for $i = 1, \ldots, N$, and $\mathbf{w}_{N+i} = [\mathbf{u}_{N-i+1}; -\mathbf{v}_{N-i+1}]$ for $i = 1, \ldots, N$, are eigenvectors of $\mathbf{C} := \begin{bmatrix} \mathbf{0} & \mathbf{T} \\ \mathbf{T}^\top & \mathbf{0} \end{bmatrix}$, with corresponding eigenvalues (in descreasing order)

$$\begin{aligned} \theta_1 = \sigma_1(\mathbf{T}), \quad \ldots, \quad \theta_N = \sigma_N(\mathbf{T}), \\ \theta_{N+1} = -\sigma_N(\mathbf{T}), \quad \ldots, \quad \theta_{2N} = -\sigma_1(\mathbf{T}). \end{aligned}$$

Define also $\tilde{\mathbf{C}} := \begin{bmatrix} \mathbf{0} & \tilde{\mathbf{T}} \\ \tilde{\mathbf{T}}^\top & \mathbf{0} \end{bmatrix}$ and its set of eigenvectors $\{\tilde{\mathbf{w}}_i\}_{i=1}^{2N}$ similarly, with $\tilde{\mathbf{w}}_1 = [\tilde{\boldsymbol{\alpha}}'; \tilde{\boldsymbol{\beta}}']$, and corresponding eigenvalues $\{\tilde{\theta}_i\}_{i=1}^{2N}$. The rest of this step is similar to Smale and Zhou [2009, Prop. 2].

Note that $\{\mathbf{w}_i\}_{i=1}^{2N}$ constitute an orthogonal basis of $\mathbb{R}^{2N}$. Thus, we can write $\tilde{\mathbf{w}}_1$ as a linear combination of this basis:

$$\tilde{\mathbf{w}}_1 = a_1 \mathbf{w}_1 + \cdots + a_{2N} \mathbf{w}_{2N},$$

with $\sum_{i=1}^{2N} a_i^2 = 1$ since all vectors in the equality have equal length of $\sqrt{2}$.

Consider $\tilde{\mathbf{C}}\tilde{\mathbf{w}}_1 - \mathbf{C}\tilde{\mathbf{w}}_1$. It can be written as

$$\tilde{\mathbf{C}}\tilde{\mathbf{w}}_1 - \mathbf{C}\tilde{\mathbf{w}}_1 = \tilde{\theta}_1 \tilde{\mathbf{w}}_1 - \sum_{i=1}^{2N} a_i \mathbf{C}\mathbf{w}_i = \sum_{i=1}^{2N} (\tilde{\theta}_1 - \theta_i) a_i \mathbf{w}_i.$$

For $i > 1$, according to condition 2, we have

$$\left| \tilde{\theta}_1 - \theta_i \right| \ge |\theta_1 - \theta_i| - \left| \theta_1 - \tilde{\theta}_1 \right| \ge |\theta_1 - \theta_2| - \left| \theta_1 - \tilde{\theta}_1 \right| \ge \frac{r}{2}.$$

As a result,

$$\left\| \tilde{\mathbf{C}}\tilde{\mathbf{w}}_1 - \mathbf{C}\tilde{\mathbf{w}}_1 \right\|^2 = 2 \sum_{i=1}^{2N} (\tilde{\theta}_1 - \theta_i)^2 a_i^2 \ge 2 \sum_{i=2}^{2N} (\tilde{\theta}_1 - \theta_i)^2 a_i^2 \ge \frac{r^2}{2} \sum_{i=2}^{2N} a_i^2. \tag{18}$$

On the other hand, we must have

$$\left\|\tilde{\mathbf{C}}\tilde{\mathbf{w}}_1 - \mathbf{C}\tilde{\mathbf{w}}_1\right\|^2 \leq \left\|\tilde{\mathbf{C}} - \mathbf{C}\right\|^2 \|\tilde{\mathbf{w}}_1\|^2 = 2\left\|\tilde{\mathbf{C}} - \mathbf{C}\right\|^2.$$

Also, $\left\|\tilde{\mathbf{C}} - \mathbf{C}\right\| = \left\|\tilde{\mathbf{T}} - \mathbf{T}\right\|$ due to the block structure of $\tilde{\mathbf{C}}$ and $\mathbf{C}$. Thus from (18) we obtain

$$\sum_{i=2}^{2N} a_i^2 \leq \frac{4}{r^2}\left\|\tilde{\mathbf{T}} - \mathbf{T}\right\|^2. \tag{19}$$

In view of the inequality $1 - x \leq \sqrt{1-x^2}$ for $x \in [0,1]$, we also have[5]

$$1 - a_1 \leq \sqrt{1 - a_1^2} = \sqrt{\sum_{i=2}^{2N} a_i^2}. \tag{20}$$

Now we can bound the perturbation

$$\begin{aligned}
\left\|\boldsymbol{\alpha}' - \tilde{\boldsymbol{\alpha}}'\right\| &\leq \|\mathbf{w}_1 - \tilde{\mathbf{w}}_1\| \\
&= \left\|(1-a_1)\mathbf{w}_1 - \sum_{i=2}^{2N} a_i \mathbf{w}_i\right\| \\
&\leq \|(1-a_1)\mathbf{w}_1\| + \left\|\sum_{i=2}^{2N} a_i \mathbf{w}_i\right\| \\
&= \sqrt{2}\left(|1-a_1| + \sqrt{\sum_{i=2}^{2N} a_i^2}\right) \\
&\leq 2\sqrt{2}\sqrt{\sum_{i=2}^{2N} a_i^2} \leq \frac{4\sqrt{2}}{r}\left\|\tilde{\mathbf{T}} - \mathbf{T}\right\|,
\end{aligned}$$

where we used that $\boldsymbol{\alpha}' - \tilde{\boldsymbol{\alpha}}'$ is a subvector of $\mathbf{w}_1 - \tilde{\mathbf{w}}_1$, the triangle inequality, (20) and (19) in the four inequalities.

Step II. We then bound the perturbation in the actual coefficients in f and $\tilde{f}$, i.e., $\|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|$. We have

$$\begin{aligned}
\frac{\|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|}{\sqrt{N}} &= \left\|(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\boldsymbol{\alpha}' - (\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\tilde{\boldsymbol{\alpha}}'\right\| \\
&\leq \left\|(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}(\boldsymbol{\alpha}' - \tilde{\boldsymbol{\alpha}}')\right\| + \left\|\left((\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1} - (\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\right)\tilde{\boldsymbol{\alpha}}'\right\| \\
&\leq \left\|(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\right\|\left\|\boldsymbol{\alpha}' - \tilde{\boldsymbol{\alpha}}'\right\| + \left\|(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1} - (\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\right\|\left\|\tilde{\boldsymbol{\alpha}}'\right\|.
\end{aligned}$$

where the triangle inequality is used in the first inequality.

We have already shown that

$$N\lambda_1\left\|(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1} - (\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\right\| = \left\|\overline{\mathbf{K}}_1(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1} - \overline{\mathbf{L}}_1(\overline{\mathbf{L}}_1 + N\lambda_1\mathbf{I})^{-1}\right\| \leq \epsilon/2$$

in Theorem 1 using the stated sampling strategy.

Using the facts $\left\|(\overline{\mathbf{K}}_1 + N\lambda_1\mathbf{I})^{-1}\right\| \leq \frac{1}{N\lambda_1}$, $\left\|\tilde{\boldsymbol{\alpha}}'\right\| = 1$, and the bound of $\left\|\boldsymbol{\alpha}' - \tilde{\boldsymbol{\alpha}}'\right\|$ from Step I, we have

$$\frac{\|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|}{\sqrt{N}} \leq \frac{1}{N\lambda_1}\frac{4\sqrt{2}}{r}\epsilon + \frac{1}{N\lambda_1}\frac{\epsilon}{2} = \left(\frac{1}{2} + \frac{4\sqrt{2}}{r}\right)\frac{\epsilon}{N\lambda_1}.$$

Step III. We now bound the kernel stability. By the Cauchy-Schwarz inequality, we have

$$\left|f(\mathbf{x}) - \tilde{f}(\mathbf{x})\right| = \left|\mathbf{k}^\top\mathbf{H}(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}})\right| \leq \|\mathbf{k}\|\,\|\mathbf{H}(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}})\| \leq \|\mathbf{k}\|\,\|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|,$$

where we have used $\|\mathbf{H}\| = 1$ in the last inequality. Since $\mathbf{k}$ is an N-dimensional vector and each entry is bounded by c (condition 1), we have $\|\mathbf{k}\| \leq \sqrt{N}c$. Combining the bound of $\|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|$ from Step II, we conclude the proof. □

[5]Without loss of generality, we assume $a_1 = \tilde{\mathbf{w}}_1^\top\mathbf{w}_1/2 \geq 0$. Otherwise, we can set $\tilde{\mathbf{w}} = [-\tilde{\boldsymbol{\alpha}}'; -\tilde{\boldsymbol{\beta}}']$ as $-\tilde{\boldsymbol{\alpha}}'$ and $-\tilde{\boldsymbol{\beta}}'$ are also a valid left/right singular vector pair of $\tilde{\mathbf{T}}$.

Remarks Theorem 2 can be extended to the case where L-dimensional projections are sought for $L > 1$, by assuming a gap at the L-th singular value of $\mathbf{T}$. Our proof technique is general and works for other low-rank KCCA algorithms, if only a bound on $\left\|\mathbf{T} - \tilde{\mathbf{T}}\right\|$ is available. With Theorem 2, one may show the consistency of NKCCA, by controlling kernel stability to be much smaller than the generalization bound of exact KCCA [Fukumizu et al., 2007].

We note in passing that, unlike kernel ridge regression [Rudi et al., 2015], KCCA does not have a representer theorem where f is restricted to the selected landmarks, as the centering operation in RKHS already involves all training samples. When the training set is large, it is worth exploring approximation strategies (see, e.g., Hsieh et al., 2014 and the references therein) to speed up prediction.

3 Related work

There is a rich literature for the Nyström method, exploring the sampling strategies and approximation quality [Drineas and Mahoney, 2005, Zhang et al., 2008, Gittens, 2011, Farahat et al., 2011, Kumar et al., 2012, Wang and Zhang, 2013, Gittens and Mahoney, 2013], applications for large-scale kernel machines [Williams and Seeger, 2000, Fowlkes et al., 2004, Platt, 2005, Zhang and Kwok, 2010, Li et al., 2010], and generalization performance of kernel machines with Nyström kernel approximations [Cortes et al., 2010, Yang et al., 2012, Jin et al., 2013, Bach, 2013, Alaoui and Mahoney, 2015, Rudi et al., 2015].

The ridge leverage scores are closely related to the *leverage scores relative to the best rank-k space* [Gittens and Mahoney, 2013], as the $\frac{\sigma_j(\mathbf{K})}{\sigma_j(\mathbf{K})+N\gamma}$ terms in (10) implement "soft shrinkage" of the spectrum of $\mathbf{K}$. In general, leverage scores measure the extent to which each sample "stands out", and are useful in a wide range of fields [Mahoney, 2011].

Avron et al. [2013] and Paul [2015] have recently proposed randomized sample selection algorithms to scale up linear CCA. Their goal is to select a subset of *paired* examples, whose canonical correlation well approximates that of the entire training set. The algorithm of Avron et al. [2013] first homogenizes the importance of each sample via a structured random projection of the training set and then uses uniform sampling on top; the algorithm of Paul [2015] uses leverage scores of certain (carefully constructed) matrix for sampling. Note that in KCCA, we are not constrained to sample data pairs; we can approximate the kernel of each view using independent columns. On the other hand, theoretically, there is little distinction between random projection and column selection: The matrices $\mathbf{S}_1$ and $\mathbf{S}_2$ can be implementing random projections instead of random column selection, as long as they possess certain spectral properties [Gittens and Mahoney, 2013].

Another approach to large-scale kernel machines is to use random Fourier features [Rahimi and Recht, 2008, 2009]. Lopez-Paz et al. [2014] map original inputs of each view to high-dimensional random feature spaces and then run linear CCA on top to approximate KCCA. While random Fourier features are data-independent and efficient to generate, the performance of this approach tends to be worse than that of the Nyström method for the same approximation rank [Lopez-Paz et al., 2014, Wang et al., 2015, Xie et al., 2015]. A key difference between the two approaches is that random Fourier features are designed to approximate the kernel functions, while the Nyström method aims to approximate the kernel matrix [Yang et al., 2012].

4 Experiments

Synthetic dataset We first verify our theoretical results on the synthetic dataset used by Fukumizu et al. [2007]. The input data are generated as follows

$$
\begin{gathered}
Z_i \sim \text{Uniform}[0,1], \\
\eta_i^x \sim \mathcal{N}(0,\, 0.02), \quad U_i = Z_i + 0.06 + \eta_i^x, \\
\eta_i^y \sim \mathcal{N}(0,\, 0.03), \quad V_i = Z_i + 3 + \eta_i^y, \\
R_i^x = \sqrt{-4\log(U_i/1.5)}, \quad \theta_i^x \sim \text{Uniform}[0, 2\pi], \\
R_i^y = \sqrt{-4\log(V_i/4.1)}, \quad \theta_i^y \sim \text{Uniform}[0, 2\pi], \\
\mathbf{x}_i = [R_i^x \cos\theta_i^x, R_i^x \sin\theta_i^x]^\top,\ \mathbf{y}_i = [R_i^y \cos\theta_i^y, R_i^y \sin\theta_i^y]^\top.
\end{gathered}
$$

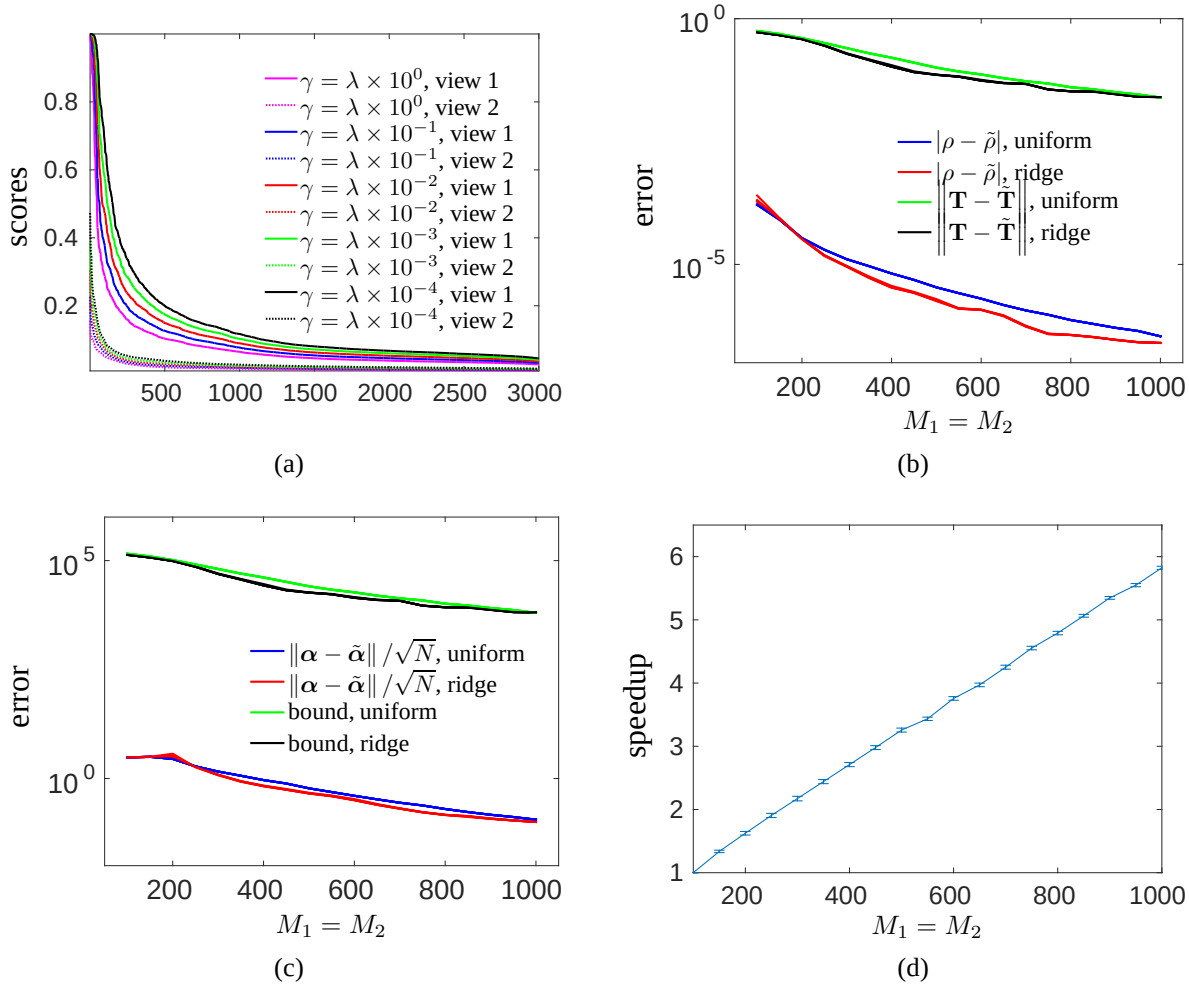

Figure 1: Results on synthetic data. (a) γ-ridge leverage scores (in sorted order) for each view. (b) $|\rho - \tilde{\rho}|$ and $\left\|\mathbf{T} - \tilde{\mathbf{T}}\right\|$. (c) $\frac{\|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|}{\sqrt{N}}$ vs. the bound $\left(\frac{1}{2} + \frac{4\sqrt{2}}{r}\right) \frac{\left\|\mathbf{T} - \tilde{\mathbf{T}}\right\|}{N\lambda_1}$. (d) Speedups achieved by our incremental algorithm at intermediate ranks. Curves in (b)(c)(d) show averaged results over 20 random seeds.

We sample $N = 3000$ points as the training set and another 3000 points as the tuning set. We perform exact KCCA and NKCCA with both uniform sampling and ridge leverage scores sampling on the training set to compute the first canonical correlation, and compare their approximation error $|\tilde{\rho} - \rho|$. Gaussian kernel widths and regularization parameters (λ_1, λ_2) are selected for best canonical correlation on the tuning set with KCCA. Exact γ-ridge leverage scores are computed at $\gamma = \lambda \times \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$ for both views and used to compute the sampling probabilities (12), and they all lead to very similar results on this dataset. We plot the γ-ridge leverage scores (in sorted order) for each view in Figure 1(a). Observe that the scores for view 1 are less uniformly distributed. We show the approximation error for a wide range of ranks ($M_1 = M_2$) from 100 to 1000 in Figure 1(b), and show in (c) the kernel stability and specifically the result of Theorem 2 step **II** (bound in step **III** can be loose), i.e., $\frac{\|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|}{\sqrt{N}}$ vs. $\left(\frac{1}{2} + \frac{4\sqrt{2}}{r}\right) \frac{\left\|\mathbf{T} - \tilde{\mathbf{T}}\right\|}{N\lambda_1}$. We see that the actual approximation errors in canonical correlation and kernel stability are smaller than the theoretical bounds, with non-uniform sampling performing somewhat better than uniform sampling.

To demonstrate the efficiency of our incremental Nyström approximation algorithm for NKCCA, we run it for up to $M_1 = M_2 = 1000$ and output the solutions at ranks $\{100, 150, 200, \ldots, 1000\}$. At each intermediate rank, We compare the run time of the incremental approach with the total run time of the non-incremental algorithm (starting fresh for each smaller rank), and plot the speedups in Figure 1(d). It is clear that the speedup improves as the number of intermediate test points increases.

Table 1: Total canonical correlations (↑) achieved by different approximate KCCA algorithms on JW11 test set with various ranks.

rank $M_1 = M_2$	1000	2000	3000	4000	5000
RCCA	69.4	82.6	88.8	92.9	95.8
NKCCA unif.	78.5	89.4	94.6	98.4	100.8
NKCCA ridge	**79.6**	**90.5**	**95.7**	**99.2**	**101.5**

Acoustic-articulatory data We next experiment with a subset of the University of Wisconsin X-Ray Microbeam corpus Westbury [1994], which consists of simultaneously recorded acoustic and articulatory measurements during speech. The acoustic view inputs are 39D mel-frequency cepstral coefficients and the articulatory view inputs are horizontal/vertical displacements of 8 pellets attached to different parts of the vocal tract, each then concatenated over a 7-frame context window. As in Lopez-Paz et al. [2014], we reshuffle all frames of speaker 'JW11', and split them into 30K/10K/10K frames for training/tuning/test. The projection dimensionality L is set to 112, which is also the maximum possible total canonical correlation.

We compare the randomized KCCA (RCCA) algorithm of Lopez-Paz et al. [2014] using random Fourier features, NKCCA with uniform sampling, and NKCCA with ridge leverage scores sampling. Gaussian kernel widths and regularization parameters are selected on the tuning set. We compute approximate γ-ridge leverage scores using the algorithm in Alaoui and Mahoney [2015, Sec. 3.5] with 5000 randomly sampled columns. γ is tuned over $\lambda \times \{10^{-2}, 10^{-1}, 1, 10^{1}\}$ and $\gamma = \lambda$ is chosen. Total canonical correlations achieved by different algorithms on the test set, averaged over 5 random seeds, are reported in Table 1 for several ranks $(M_1 = M_2)$. We observe that NKCCAs outperform RCCA with the same approximation rank, which is consistent with previous work, and sampling based on approximate ridge leverage scores improves over uniform sampling by a small margin.

Discussion Empirically, we find that non-uniform sampling based on ridge leverage scores provides moderate improvement over uniform sampling. Several factors are at play. First, ridge leverage scores are data-dependent quantities, and theoretically there is a clear advantage to use non-uniform sampling if the ridge leverage scores are quite non-uniform. Second, for large datasets, it is computationally infeasible to compute exact ridge leverage scores, and approximating them results in sub-optimal performance. Third, although ridge leverage scores provide better approximation for kernel ridge regression where a single kernel is involved, the interplay between the two kernels is also crucial in KCCA. This interplay is not fully taken into account by our analysis (e.g., the upper bound in (14) consists of error terms for each view individually) and our sampling strategy. It is an interesting future direction to study sampling strategies that use information from the other view, which may lead to improved approximation accuracy.

5 Conclusions

We have proposed a non-uniform sampling strategy for column selection in approximate KCCA using the Nyström method, proved its approximation guarantees for both training (error in canonical correlation) and testing (kernel stability), and also developed an algorithm for computing solutions incrementally at different approximation ranks.

It has been shown by Bach [2013], Alaoui and Mahoney [2015] that for kernel ridge regression, which has a closed-form solution (as does KCCA), it is possible to derive a sharper generalization bound directly using the solution with low-rank kernel approximations, rather than indirectly through the kernel stability [Cortes et al., 2010]. Therefore, another interesting problem is to develop a sharp statistical performance guarantee of NKCCA based on known results for exact KCCA [Fukumizu et al., 2007].

A Proof of Lemma 2

Proof. We first prove the bound for uncentered kernel matrices. Notice that $\mathbf{K}(\mathbf{K}+N\lambda\mathbf{I})^{-1} = \mathbf{I} - N\lambda(\mathbf{K}+N\lambda\mathbf{I})^{-1}$. Therefore, for any $\gamma \geq 0$,

$$\begin{aligned}
&\left\|\mathbf{K}(\mathbf{K}+N\lambda\mathbf{I})^{-1} - \mathbf{L}_\gamma(\mathbf{L}_\gamma+N\lambda\mathbf{I})^{-1}\right\| \\
&= \left\|\left(\mathbf{I} - N\lambda(\mathbf{K}+N\lambda\mathbf{I})^{-1}\right) - \left(\mathbf{I} - N\lambda(\mathbf{L}_\gamma+N\lambda\mathbf{I})^{-1}\right)\right\| \\
&= N\lambda\left\|(\mathbf{K}+N\lambda\mathbf{I})^{-1} - (\mathbf{L}_\gamma+N\lambda\mathbf{I})^{-1}\right\| \\
&= N\lambda\left\|(\mathbf{K}+N\lambda\mathbf{I})^{-1}(\mathbf{L}_\gamma-\mathbf{K})(\mathbf{L}_\gamma+N\lambda\mathbf{I})^{-1}\right\| \\
&\leq N\lambda\left\|(\mathbf{K}+N\lambda\mathbf{I})^{-1}\right\|\|\mathbf{K}-\mathbf{L}_\gamma\|\left\|(\mathbf{L}_\gamma+N\lambda\mathbf{I})^{-1}\right\| \\
&\leq \frac{1}{N\lambda}\|\mathbf{K}-\mathbf{L}_\gamma\| \qquad (21)
\end{aligned}$$

where we have used the fact $\left\|(\mathbf{K}+N\lambda\mathbf{I})^{-1}\right\| \leq \frac{1}{N\lambda}$ in the last inequality.

For $\gamma > 0$, we can use $\mathbf{K}-\mathbf{L}_\gamma \preceq \frac{N\gamma}{1-t}\mathbf{I}$ from Lemma 1 to obtain the desired bound. On the other hand, we observe that $\mathbf{K}-\mathbf{L} \preceq \mathbf{K}-\mathbf{L}_\gamma$ for any $\gamma > 0$ owing to $\mathbf{L}_\gamma \leq \mathbf{L}$ from Lemma 1. Therefore $\|\mathbf{K}-\mathbf{L}\| \leq \|\mathbf{K}-\mathbf{L}_\gamma\|$ and the same bound holds for $\mathbf{L}$.

Now, for centered matrices, a derivation similar to that of (21) follows and we just need to bound $\left\|\overline{\mathbf{K}}-\overline{\mathbf{L}}_\gamma\right\|$. By definitions of the centered matrices and conjugating the PSD ordering [Horn and Johnson, 1986, Obs. 7.7.2], we have

$$\overline{\mathbf{K}}-\overline{\mathbf{L}} \preceq \overline{\mathbf{K}}-\overline{\mathbf{L}}_\gamma \preceq \frac{N\gamma}{1-t}\mathbf{H}^2 \preceq \frac{N\gamma}{1-t}\mathbf{I} \qquad (22)$$

where the last step is due to the fact that eigenvalues of $\mathbf{H}$ are in $\{0,1\}$. Thus the reasoning for uncentered matrices also holds here. □

References

S. Akaho. A kernel method for canonical correlation analysis. In *Proceedings of the International Meeting of the Psychometric Society (IMPS2001)*, 2001.

A. E. Alaoui and M. W. Mahoney. Fast randomized kernel methods with statistical gaurantee. In *NIPS*, pages 775–783, 2015.

R. Arora and K. Livescu. Kernel CCA for multi-view learning of acoustic features using articulatory measurements. In *Symposium on Machine Learning in Speech and Language Processing (MLSLP)*, 2012.

H. Avron, C. Boutsidis, S. Toledo, and A. Zouzias. Efficient dimensionality reduction for canonical correlation analysis. In *ICML*, pages 347–355, 2013.

F. Bach. Sharp analysis of low-rank kernel matrix approximations. In *Proc. of the 26th Annual Conference on Learning Theory (COLT'13)*, pages 1–25, 2013.

F. R. Bach and M. I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.

M. B. Blaschko, J. A. Sheltonb, A. Bartelsc, C. H. Lamperte, and A. Gretton. Semi-supervised kernel canonical correlation analysis with application to human fMRI. *Pattern Recognition Letters*, 32(11):1572–1583, 2011.

C. Cortes, M. Mohri, and A. Talwalkar. On the impact of kernel approximation on learning accuracy. In *Proc. of the 13th Int. Workshop on Artificial Intelligence and Statistics (AISTATS 2010)*, pages 113–120, 2010.

P. Drineas and M. W. Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6:2153–2175, 2005.

P. Drineas, R. Kannan, and M. W. Mahoney. Fast Monte Carlo algorithms for matrices I: Approximating matrix multiplication. *SIAM J. Comp.*, 36(1):132–157, 2006.

A. Farahat, A. Ghodsi, and M. Kamel. A novel greedy algorithm for Nyström approximation. In *Proc. of the 14th Int. Workshop on Artificial Intelligence and Statistics (AISTATS 2011)*, pages 269–277, 2011.

C. Fowlkes, S. Belongie, F. Chung, and J. Malik. Spectral grouping using the Nyström method. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 26(2):214–225, 2004.

K. Fukumizu, F. R. Bach, and A. Gretton. Statistical consistency of kernel canonical correlation analysis. *Journal of Machine Learning Research*, 8:361–383, 2007.

A. Gittens. The spectral norm error of the naive Nystrom extension. arXiv:1110.5305 [math.NA], 2011.

A. Gittens and M. Mahoney. Revisiting the nystrom method for improved large-scale machine learning. In *ICML*, pages 567–575, 2013.

G. H. Golub and C. F. van Loan. *Matrix Computations*. third edition, 1996.

D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural Computation*, 16(12):2639–2664, 2004.

D. R. Hardoon, J. Mourão-Miranda, M. Brammer, and J. Shawe-Taylor. Unsupervised analysis of fMRI data using kernel canonical correlation. *NeuroImage*, 37(4):1250–1259, 2007.

M. Hodosh, P. Young, and J. Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899, 2013.

R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1986.

H. Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.

C.-J. Hsieh, S. Si, and I. S. Dhillon. Fast prediction for large-scale kernel machines. In *NIPS*, pages 3689–3697, 2014.

R. Jin, T. Yang, M. Mahdavi, Y.-F. Li, and Z.-H. Zhou. Improved bounds for the Nyström method with application to kernel classification. *IEEE Trans. Information Theory*, 59(10):6939–6949, 2013.

S. Kumar, M. Mohri, and A. Talwalkar. Sampling methods for the Nyström method. *Journal of Machine Learning Research*, 13:981–1006, 2012.

P. L. Lai and C. Fyfe. Kernel and nonlinear canonical correlation analysis. *Int. J. Neural Syst.*, 10(5):365–377, 2000.

M. Li, J. T. Kwok, and B. Lu. Making large-scale Nyström approximation possible. In *ICML*, pages 631–638, 2010.

Y. Li and J. Shawe-Taylor. Using KCCA for japanese-english cross-language information retrieval and classification. *Journal of Intelligent Information Systems*, 27(2):117–133, 2005.

D. Lopez-Paz, P. Hennig, and B. Schölkopf. The randomized dependence coefficient. In *NIPS*, pages 1–9, 2013.

D. Lopez-Paz, S. Sra, A. Smola, Z. Ghahramani, and B. Schoelkopf. Randomized nonlinear component analysis. In *ICML*, pages 1359–1367, 2014.

M. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2): 123–224, 2011.

T. Melzer, M. Reiter, and H. Bischof. Nonlinear feature extraction using generalized canonical correlation analysis. In *Proc. of the 11th Int. Conf. Artificial Neural Networks (ICANN'01)*, pages 353–360, 2001.

S. Paul. Core-sets for canonical correlation analysi. In *Proc. of the 24nd Int. Conf. Information and Knowledge Management (CIKM'15)*, 2015.

J. Platt. FastMap, MetricMap, and landmark MDS are all Nyström algorithms. In *Proc. of the 10th Int. Workshop on Artificial Intelligence and Statistics (AISTATS 2005)*, pages 261–268, 2005.

A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, pages 1177–1184, 2008.

A. Rahimi and B. Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *NIPS*, pages 1313–1320, 2009.

A. Rudi, R. Camoriano, and L. Rosasco. Less is more: Nyström computational regularization. In *NIPS*, pages 1648–1656, 2015.

F. Rudzicz. Correcting errors in speech recognition with articulatory dynamics. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 60–68, 2010.

S. Smale and D.-X. Zhou. Geometry on probability spaces. *Constructive Approximation*, 30(311–323), 2009.

R. Socher and F.-F. Li. Connecting modalities: Semi-supervised segmentation and annotation of images using unaligned text corpora. In *Proc. of the 2010 IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR'10)*, pages 966–973, 2010.

A. Vinokourov, N. Cristianini, and J. Shawe-Taylor. Inferring a semantic representation of text via cross-language correlation analysis. In *NIPS*, pages 1497–1504, 2003.

S. Wang and Z. Zhang. Improving CUR matrix decomposition and the Nyström approximation via adaptive sampling. *Journal of Machine Learning Research*, 14:2729–2769, 2013.

W. Wang, R. Arora, K. Livescu, and J. Bilmes. On deep multi-view representation learning. In *ICML*, pages 1083–1092, 2015.

J. R. Westbury. *X-Ray Microbeam Speech Production Database User's Handbook Version 1.0*. University of Wisconsin, Madison, 1994.

C. K. I. Williams and M. Seeger. The effect of the input density distribution on kernel-based classifiers. In *ICML*, pages 1159–1166, 2000.

C. K. I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *NIPS*, pages 682–688, 2001.

B. Xie, Y. Liang, and L. Song. Scale up nonlinear component analysis with doubly stochastic gradients. arXiv:1504.03655 [cs.LG], 2015.

Y. Yamanishi, J.-P. Vert, and M. Kanehisa. *Kernel Methods in Computational Biology*, chapter Heterogeneous Data Comparison and Gene Selection with Kernel Canonical Correlation Analysis, pages 209–229. 2004.

T. Yang, Y.-F. Li, M. Mahdavi, R. Jin, and Z.-H. Zhou. Nyström method vs random Fourier features: A theoretical and empirical comparison. In *NIPS*, pages 476–484, 2012.

K. Zhang and J. T. Kwok. Clustered Nyström method for large scale manifold learning and dimension reduction. *IEEE Trans. Neural Networks*, 21(10):1576–1587, 2010.

K. Zhang, I. W. Tsang, and J. T. Kwok. Improved Nyström low-rank approximation and error analysis. In *ICML*, pages 1232–1239, 2008.